

Amendments to the Claims

Kindly amend claim 12, and cancel claim 20 (without prejudice), as set forth below. All pending claims are reproduced below, with changes in the amended claims shown by underlining (for added matter) and strikethrough/double brackets (for deleted matter).

1-11. (Previously Canceled).

12. (Currently Amended) A method of processing comprising:

providing, by a dedicated collective offload engine coupled to a switch fabric in a distributed, parallel computing system, collective processing of data from at least some processing nodes of multiple processing nodes of the distributed, parallel computing system, the dedicated collective offload engine being a hardware device coupled to the switch fabric, the hardware device being a specialized device dedicated to providing collective processing ~~in hardware~~ of data from the at least some processing nodes and comprising a dispatcher built from field programmable gate arrays, and a pipelined arithmetic logic unit, the dispatcher controlling collective processing of data by the arithmetic logic unit, the collective processing implementing a collective operation on the data from the at least some processing nodes without use of a software tree;

producing, by the dedicated collective offload engine, a result based on said collective processing; and

forwarding said result to at least one processing node of the multiple processing nodes.

13. (Previously Presented) The method of claim 12, wherein the collective operation is a Message Passing Interface (MPI) collective operation.

14. (Original) The method of claim 12, further comprising:

receiving and storing, at a payload memory, the data from the at least some processing nodes of the multiple processing nodes, wherein said payload memory is a component of the dedicated collective offload engine; and

retrieving and performing, at an arithmetic logic unit (ALU), the collective processing of data stored in the payload memory, wherein said ALU is a component of the dedicated collective offload engine and is coupled to the payload memory.

15. (Original) The method of claim 14, further comprising:

controlling the collective processing of the data from the at least some processing nodes of the multiple processing nodes, wherein said controlling is performed by a dispatcher of the dedicated collective offload engine coupled to the ALU, and in communication with the at least some processing nodes of the multiple processing nodes via the switch fabric; and

controlling, by the dispatcher, the sharing of the result with the at least one processing node of the multiple processing nodes.

16. (Original) The method of claim 15, further comprising:

storing, in at least one task table coupled to the dispatcher, task identification information related to the at least some processing nodes of the multiple processing nodes, wherein said at least one task table is a component of the dedicated collective offload engine; and

storing, in at least one synchronization group table coupled to the dispatcher, identification information related to one or more groups of the at least some processing nodes of the multiple processing nodes, wherein said at least one synchronization group table is a component of the dedicated collective offload engine.

17. (Original) The method of claim 15, further comprising:

communicating, via an adapter, across the switch fabric using a link protocol, wherein said adapter is coupled to the switch fabric and is a component of the dedicated collective offload engine; and

facilitating, by interface logic, communication between said adapter and said payload memory and between said adapter and said dispatcher, wherein said interface logic is a component of the dedicated collective offload engine.

18. (Original) The method of claim 12, further comprising:

communicating among a plurality of dedicated collective offload engines via the switch fabric, wherein said communicating facilitates the collective processing of data from the at least some processing nodes of the multiple processing nodes and the producing of the result based thereon.

19. (Original) The method of claim 12, further comprising:

communicating among a plurality of dedicated collective offload engines via a channel disposed therebetween, said channel being independent of the switch fabric, wherein said communicating facilitates the collective processing of data from the at least some processing nodes of the multiple processing nodes and the producing of the result based thereon.

20. (Canceled).

21-30. (Previously Canceled).

* * * * *